Utilizing IoT Devices for Monitoring and Adjusting Clinical Pathway Exercises

Nicole Merkle[1] and Stefan Zander[2]

[1] FZI Forschungszentrum Informatik am KIT, Information Process Engineering, Haid-und-Neu-Str. 10-14, D-76131 Karlsruhe, Germany, merkle@fzi.de
[2] University of Applied Sciences Darmstadt, Institute for Computer Science, Schöfferstrasse 8B, D-64295 Darmstadt, Germany, stefan.zander@h-da.de

Abstract. Clinical pathways play a crucial role in the rehabilitation process of patients during and after a clinical stay. Posterior to the clinical stay, patients are usually left to their own devices. This fact hampers the adequate treatment of the patient as there is no longer a technical controlling and guiding instance. For this reason, this work describes the current status of the virtual coach, a) a semantic representation framework for clinical pathways with their immanent exercises and combines this semantic representation framework with b) a reinforcement learning approach for computing the optimal exercise severity level according to the current patient capabilities and recovery. The semantic representation and linkage of heterogeneous Internet of Things (IoT) devices to medical exercises helps to monitor and adjust exercise recommendations according to sensed environmental states. The combination of IoT with semantic web technologies as well as reinforcement learning contributes towards an optimal rehabilitation process.

Keywords: Agent-Based System, Semantic Web, Healthcare, Context-Awareness, Reinforcement Learning, IoT, WoT

1 Introduction

Patients who have lost or weakened their motoric capabilities due to an accident or disease have to perform—during as well as after their medical treatment—exercises in order to train their lost motoric capabilities. The exercises are defined by the patient's physician or physiotherapist and follow a strict and concrete sequence. Usually, the patient's physiotherapist has to coach and check during the exercises if they are performed correctly by the patient. This requires that the patient has to perform these exercises during a hospital stay in frequent sessions under control of the physiotherapist. Mostly, after the hospital stay, the exercises have—depending on the disease or impairment—to be continued by the patient alone. However, the patient has no instructor who controls and supports the patient during the exercises repetition what can lead to a wrong appliance of exercises and worsen the recovery progress. This observation shows the necessity of a virtual coaching system for coaching the patient individually, in particular after the hospital stay through appropriate and adequate exercises.

This works aims at addressing certain IoT challenges using semantic web technologies. Especially, monitoring and making sense out of sensed data together with adaptability are crucial challenges that need to be addressed in medical recovery scenarios. Moreover, we describe a concept for semantically representing clinical pathway exercises in relation with the IoT related observation system. Based on semantic exercise descriptions, reinforcement learning (RL) is applied in order to compute and create semantic optimal exercise execution paths (policies). By means of this execution paths, the virtual coach is enabled to check the steps of the performed tasks and to recommend the next exercise steps. Moreover, the virtual coaching system can notice by the exercise progress which capability has been trained successfully in order to inform the patient about his/her trained capabilities.

The remainder of this work is structured as follows. Section 2 discusses related work in order to distinguish this work from other approaches. Section 3 describes a possible use case scenario of a patient and his recovery process. Section 4 presents the semantic representation of training exercises as well as required capabilities and shows how the computation of optimal task executions by means of RL is performed. Section 5 illustrates the proof-of-concept of the presented approach. Section 6 ends with a conclusion and gives an outlook to future work.

2 Related Work

A number of frameworks for monitoring and adjusting clinical pathways to the patient's characteristics already exist; however, only few of them incorporate semantic technologies and almost none of them considers IoT technology and the possibility of utilizing such devices for rehabilitation purposes. In the following, we introduce the most relevant related works.

Dragoni et. al. [3] propose a semantic platform for monitoring patient activities. Their proposed platform recommends—by reasoning rules—healthy activities (e.g. healthy diets), based on the user generated data and the presented domain knowledge. Bailoni et al. [2] developed *PerKApp*, a context-aware and user-centric platform utilizing "persuasion technologies, natural language and deep knowledge representation tools" for proposing healthier lifestyles to target users. Wang et al. [10] propose a "generic framework for the hospital-specific customization of standard care plans defined by clinical pathways or clinical guidelines". Therefore, it constructs a semantic data model in order to store semantic clinical data and extract ordered treatment procedures. An ontology model for generating clinical pathways by adopting organizational semiotics has been developed by Tehrani et al. [9]. The authors apply the Semantic Analysis Method (SAM) in order to represent clinical pathways and the Norm Analysis Method (NAM) for recognizing behavioral patterns and rules of clinical pathways. The approach by Li et al. [6] presents a knowledge representation framework which is built upon inputs of organizational semiotics. This computerized knowledge representation is intended to improve treatment processes as well as the quality of medical services. In 2014, Li et al. [5] Li developed a norm-based approach for managing clinical pathways and upon these simulating a multi-agent system. In this approach, norms are represented by rules. Their approach aims at the integration of pathway knowledge into treatment processes and hospital information systems. The personalization and adaption of healthcare processes and treatment plans was also discussed by Alexandrou et al. [1]. The authors utilize an ontology together with a set of semantic rules in order to achieve a predetermined form of process adaptability. Laleci et al. [4] developed "a clinical decision support system for remote monitoring of patients at their homes". Their main objective is to provide a semantic multi-agent-based interoperability framework for heterogeneous clinical systems. Shen et al. [7] combined a case-based reasoning approach with an ontological multi-agent architecture for facilitating clinical decision support systems. By applying case-based reasoning, solutions for known problems can be applied for solving similar problems.

All considered approaches aim at the representation and integration of clinical pathways into electronic and heterogeneous systems. To support this, ontologies, rule- and multi-agent-based systems are utilized, mainly to improve decision support systems. However, all of these approaches do neither consider the continuation of rehabilitation processes at home nor consider the potential of IoT devices. The patient as well as his/her capabilities and recovery progress are not in the primary focus of the discussed approaches. Moreover, in none of the mentioned approaches RL is applied for training agents regarding optimal clinical pathway exercises.

3 Use Case

In this section, we provide an example use case which motivates and reinforces the necessity of a virtual coaching system for processing environmental observations by means of IoT devices. Moreover, the use case involves a persona (Magnus Mirks) which represents a potential target user. The use case demonstrates general requirements, which were identified during interviews and workshops with medical domain experts (e.g. physicians, physiotherapists, caregivers).

Magnus Mirks—a 68 year old retired teacher—got an on-pump coronary artery bypass. After the surgery, he had to stay for two weeks at the hospital and a physiotherapist helped him frequently to train the capacity of his lung as well as his condition, endurance, mobility and body power. Magnus is now at home again and knows that he has to continue to train these capabilities. Every of the learned exercises are aligned to his health status and the treatment objectives which are to train the mentioned capabilities. Therefore, he needs an infrastructure for observing his exercise activities in order to evaluate his progress in training his capabilities. Considering his situation, Magnus thinks that it would be fine to have a virtual motivating game where he can perform his training exercises and achieve rewards and the next level for successfully performing his exercises. He believes that this would increase his motivation to follow the prescribed exercises and to fasten his recovery. In addition, he thinks, an avatar—interacting with him in a human like way—could coach him during the performance of exercise tasks. The example of Magnus shows the following observations: a) A framework is necessary for observing and providing his exercise executions in a processable format. b) A reward mechanism is appreciated by the patient as it increases the motivation to continue the exercises. c) A human-like interaction is required by the patient in order to get feedback regarding the performed exercises and to assure the patient about the trained capabilities. In the next section, we discuss how these requirements are addressed by the presented approach.

4 Approach

The presented virtual coaching system comprises several different system components. However, we just present three of them: a Web of Things (WoT) server for abstracting IoT devices, a RL agent for computing optimal exercise policies and a Semantic MediaWiki(SMW)[3]—enhanced with an extension and templates—for creating a semantic representation of clinical pathway relevant entities (e.g. IoT devices, user profiles, disease profiles, medical exercises). These components are discussed in the following subsections. Section 4.1 introduces the semantic representation of exercises and user capabilities. Section 4.2 shows how the semantic representation conduces to the RL approach for recognizing exercise states and computing according to these states the optimal task executions. Section 4.3 presents the exercise modelling tool[4] for physicians. This tool allows the physicians to provide semantic exercise descriptions without being an expert in semantic technologies. In order to provide a framework for physicians, we have chosen SMW. SMW is a content management platform which enables its users to provide semantic RDF(S) representations of wiki pages. These semantic representations are used by the coaching system in order to process relevant information about the patient and the medical exercises. Furthermore, semantic web technologies are applied because of their expressiveness and immanent shared understanding. They allow an adequate description and linkage of the related exercises to required capabilities and the patient's performed actions. In addition, we consider in the representation the involved sensors in order to describe exercise states and recognize the current state of performed exercise tasks. The linkage between all these semantic entities allows the virtual coaching system to compute optimal exercise paths and to provide these to the coaching avatar for supporting the patient during the exercise execution. Figure 1 illustrates the big picture of the virtual coaching framework.

[3] For more details, see: `https://www.semantic-mediawiki.org/wiki/Semantic_MediaWiki`

[4] SemanticStateChart extension for SMW

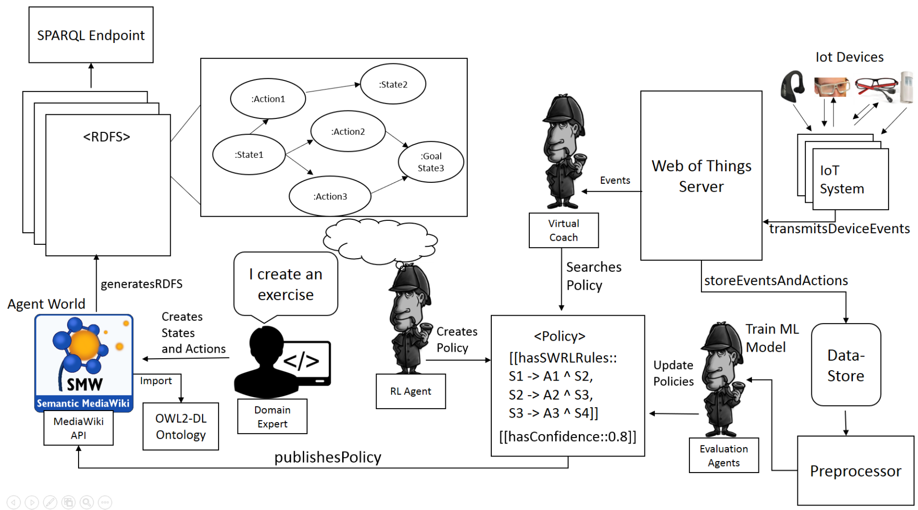

Fig. 1. The architecture of the virtual coach

The *WoT* server provides by different *IoT adapter* components an abstraction layer between heterogeneous IoT devices of the environment and the running virtual coach. The IoT device specific protocols are implemented and transferred by these IoT adapters into a general JSON-LD representation. The WoT itself uses the JSON-LD messages for communicating with the virtual coach and a data store. The data store collects sensor events and performed actions in order to provide for machine learning agents data samples. Every IoT device is described by events (evoked by the device), actions (triggering the device) and additional properties. The *virtual coach* agent subscribes at the WoT for all registered IoT devices and is informed as soon as a device is changing its state. Subsequently, the virtual coaching agent requests from the *SMW* all available state representations and deduces by them and the received IoT device states the appropriate state of the environment. A state is a representation of the environment and can be considered as an observation about the current sensed context (e.g. User is in LivingRoom, User is awake, Kitchen is dark). Moreover, every state is related to possible actions and a reward value that allows the virtual coaching agent to asses the state. In a next step, the virtual coaching agent searches by SPARQL queries, the appropriate exercise policy representation, which is linked to the given state. If a matching policy is found, the virtual coaching agent executes the suggested action and its possible sub actions. The execution of an action leads again to a new state change which is reported by the WoT server to the virtual coaching agent. Iteratively, the virtual coaching agent performs these process steps by means of the WoT and the policies until a goal state is achieved. However, before the virtual coaching agent is able to follow these steps, it has to compute beforehand the policies for given exercises. The exercises are created by a domain expert (e.g. physiotherapist, physician) via the *Semantic StateChart* extension, which is integrated in the SMW. In order to avoid mistakes by the domain expert, the SemanticStateChart extension purports some restrictions, which have to be followed by the domain expert. For instance, every state requires to be linked to an action. Even the goal state has to reference itself by an action. Moreover, every state has to provide a reward value which can be either a value of -1, 0, or 1. Every state and action has to be unique because semantic entities are unique. The mentioned model creation rules are documented in the SemanticStateChart extension in order to support the domain expert during the creation of medical exercises.

The virtual coaching agent requests exercise representations which are marked as *open* and computes for these exercises and for every related state the appropriate action (policy). After the computation, the agent sets the exercise status to *done* and creates a semantic policy representation, which it publishes in SMW, so that it can be accessed later by the virtual coaching agent in the local network. By means of these policies, the virtual coaching agent recommends to the patient the appropriate exercise actions and deduces after every successfully accomplished exercise, the trained capabilities.

4.1 Semantic Representation of Exercises and Capabilities

Every training exercise has a degree of difficulty. This degree of difficulty depends on the achievements of the patient during an exercise. If an exercise can be accomplished successfully, the degree of difficulty might be increased in order to challenge and train the capabilities of the patient. However, it depends heavily on the full execution of the exercise tasks whether a patient can achieve the next exercise level. Therefore, it is required that the virtual coaching system can recognize by e.g. optical, kinect or vital sign sensors in which execution step the patient resides. This implies that the virtual coaching system contains an expressive representation of the clinical pathway as well as a concrete description of the exercise tasks and of the monitoring IoT devices.

Moreover, trained capabilities—described by OWL 2 DL and their entities—are linked semantically to the exercise representations. The virtual coach is enabled by this to deduce by performed exercises the appropriate trained capabilities.

Furthermore, the interaction of the coaching system with a patient is necessary, in order to demonstrate or explain the moves in cases feedback is demanded by the patient. It is necessary to tightly connect the observation system with the semantic processing infrastructure in order to make sense out of acquired sensorial data—in particular to assure that every step has been performed. This section shows how sensorial data from IoT devices helps to monitor and deduce the patient's context (e.g. activities) in order to recommend appropriate actions for accomplishing the given medical exercises. The proposed semantic representation uses state-action diagrams in order to present clinical pathway exercises. The states can be considered as an aggregation of sensed sensor states and are described by means of the Semantic Web Rule language (SWRL) since SWRL rules can express complex circumstances. The rules define, if the conditions for a state are achieved. Therefore, the virtual coach observes during the exercises the appropriate sensor states and reasons by these rules whether the sensor states match to the appropriate state. A state itself represents an achieved task step. Given a current state, the virtual coach can decide if the goal state of the exercise has been already achieved by checking the semantic state representation. The definition of a goal state is prescribed by the physician during the exercise specification. Is this the case, the virtual coach knows the patient's next optimal action to perform in order to achieve the next exercise step and degree of difficulty because the exercises are linked to each other, ordered according to their degree of difficulty. It is sufficient to describe the state-action diagrams in a deterministic way, because the exercises are also deterministic and have always an absorbing goal state. The state transitions are represented by actions which lead to a next state. In order to transfer a state-action representation into a semantic one, the state diagram is transferred into a RDF(S) representation. This step is necessary, because the RDF(S) representation provides additional knowledge about linked required capabilities for every exercise, reward values for every state, compound device states, which represent an exercise state and about the processing status of a task. This knowledge is used by the virtual coach in order to reason trained capabilities. As previously mentioned, every state entity is described by SWRL rules. Equation 1 shows the general structure of such a rule:

$$\begin{aligned} &\mathsf{Patient}(?\mathsf{p}) \sqcap \mathsf{Sensor}(?\mathsf{s}) \sqcap \mathsf{hasState}(?\mathsf{s}, \text{"}\mathsf{value}\text{"}) \\ &\qquad\qquad\qquad \Rightarrow \mathsf{isInState}(?\mathsf{p}, \mathsf{State_n}) \end{aligned} \tag{1}$$

The rule expresses that a patient $?p$ is in some $State_n$, if a sensor $?s$ has sensed a certain value. The rule can contain an arbitrary count of predefined sensor states in order to describe a semantic state entity. In this way, the virtual coach observes certain sensor states in order to decide, which exercise state the patient has already achieved. The exercise state entities are linked via RDF(S) properties to possible actions, which the patient has to perform to achieve the next state. The semantic action entities describe a single movement and contain for an explanation a human readable description. A virtual coach avatar is then enabled to request—via a SPARQL endpoint—and read this description in order to support the patient in performing the current action. In this way, every exercise task is specified by a semantic graph representation, which is processed subsequently by the virtual coach.

It is important to know that every exercise sequence requires a certain policy, which the patient has to follow. The policy prescribes an action to perform in a certain exercise state. This is especially necessary if a state allows more than one actions to perform. Equation 2 shows a policy representation. However, the policy is not predefined in the exercise representation. The virtual coaching agent computes the optimal policies based on the underlying semantic exercise representation in order to provide optimal exercise paths.

$$\pi^*(\mathsf{s_t}) = a_t \tag{2}$$

Every exercise step requires a certain motoric capability in order to achieve the final state of a successfully performed exercise. In order to represent this, capability entities are linked to exercise state entities. Axiom 3 illustrates the linkage between an exercise state and a capability via a subsumed role restriction axiom in OWL 2 DL.

$$\mathsf{State} \sqsubseteq \exists\ \mathsf{requiresCapability.SomeCapability} \tag{3}$$

Using the previous axiom, the following RDF Turtle[5] representation shows how an assertional example representation of an exercise state looks like:

```
<http://192.168.2.104/vc/LegSweeping>
  rdf:type :State;
  :trainsCapability :Mobility,
                    :Coordination.
```

The exercise state entity *LegSweeping* trains two different capabilities (*Mobility, Coordination*). The virtual coaching agent can deduce implicitly through the accomplished state whether the patient provides the required capabilities for performing that exercise task. The assertional knowledge is represented by means of SMW. Every exercise entity as well as capability requires to be created and represented by a Wiki page. SMW generates by the page representation a RDF(S) representation, which can be requested via a SPARQL endpoint by the virtual coaching agent in order to compute optimal exercise paths.

4.2 The Computation of Optimal Exercise Paths

For the computation of optimal actions in certain exercise states, the virtual coaching agent applies RL, because it allows the agent to simulate and learn by a reward mechanism the optimal actions to perform in order to master an exercise. Moreover, RL does not require like supervised learning, a huge amount of datasets. Only the semantic state-action representations of exercises are sufficient for computing optimal exercise policies. A policy is an optimal strategy to achieve an objective. In this case the objective is to perform an exercise successfully in order to train lacking motoric capabilities. According to Sutton and Barto [8], a policy can be computed by a *Q-Learning function*. The function is illustrated in Equation 4.

$$\mathsf{Q}^*(\,\mathsf{S}_t, \mathsf{A}_t) = \sum_{t=0}^{T} \mathsf{r}_{t+1} + \alpha * max(\,\mathsf{Q}^*(\,\mathsf{S}_{t+1}, \mathsf{A}_{t+1})\,) \qquad 0 \le \alpha \le 1 \tag{4}$$

[5] `https://www.w3.org/TR/turtle/`

The Q-function computes for every state-action pair a cumulative reward value. This value is determined in several episodes. An episode starts at the exercise beginning and ends if the exercise goal state is achieved. As soon as all Q-function values are converging, the optimal policies are found and the virtual coach stops its computation and publishes them in a semantic representation via SMW. Hence, the found policies are subsequently transferred into SWRL rules. These SWRL rules are linked via semantic annotations to a policy entity. Axiom 5 illustrates a general rule structure of every exercise policy.

$$\begin{aligned}\mathsf{Patient}(?\mathsf{p}) \sqcap \mathsf{State}(\mathsf{State_x}) \\ \sqcap \mathsf{Action}(\mathsf{Action_x}) \sqcap \mathsf{isInState}(?\mathsf{p}, \mathsf{State_x}) \\ \Rightarrow \mathsf{hasOptimalAction}(\mathsf{State_x}, \mathsf{Action_x})\end{aligned} \tag{5}$$

The given rule defines in its premise patient, state and action entities and claims that a patient has to be in a certain $State_x$, which has an optimal action. This optimal action has been computed previously by means of the RL approach. In order to provide the computed policy, the virtual coach creates autonomously via the MediaWiki API[6] a semantically enhanced policy representation. The policy page is linked via semantic annotations to required capabilities and the exercise states, including the goal state. Therefore, the agent requests via SPARQL the appropriate matching capabilities. An example query for capability-retrieval is depicted in equation 6. The SPARQL query requests for all capabilities which are linked to an entity of the *JointRotationToStretching* action class. If more than one capabilities are returned, then the virtual coach can show the patient that all of these capabilities are trained by performing the recommended action.

$$\begin{aligned}&\mathsf{SELECT\ \ ?capability\ \ WHERE}\{\\ &\mathsf{?capability\ \ rdf:type\ \ Category:Capability.}\\ &\mathsf{?capability\ \ :hasAction\ \ ?action.}\\ &\mathsf{?action\ \ rdf:type\ \ :JointRotationToStretching.}\}\end{aligned} \tag{6}$$

The published policy is used afterwards during the exercise execution by the virtual coach for suggesting and evaluating the right exercise execution path.

4.3 The Exercise Modelling Tool

In order to support physicians at the provision of exercise descriptions, the virtual coach provides via SMW an extension (*SemanticStateChart*) for creating semantic exercise descriptions. Figure 2 illustrates the UI of the extension. The physician requires to provide a state-action diagram, an exercise task name, a goal state and a discount factor for the Q-learning function. The nodes of the state-action diagram are representing the exercise states, while the arrows are representing the possible actions. Additionally, to every node a reward value has to be assigned after a colon. The extension generates after the storage of

[6] $https://www.mediawiki.org/wiki/API:Main\_page$

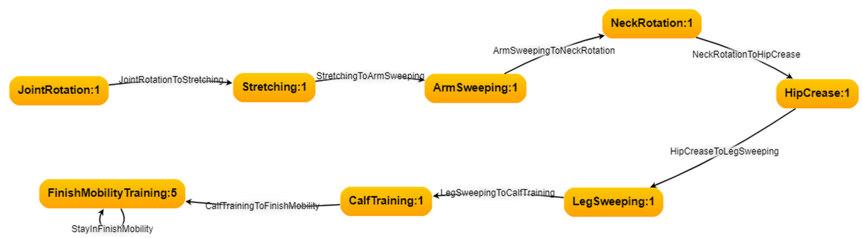

Fig. 2. The SemanticStateChart extension for a graphical exercise diagram creation

this specification, a new semantic exercise representation by creating and linking new state- and action entities in terms of annotated wiki pages. Afterwards, the physician has the possibility to enhance via forms and templates the created state entities with additional information (e.g. state-conditions as presented in Section 4.1).

5 Proof-of-Concept

For evaluation purposes, we generate a sequence of medical exercises e.g. *Train-Coordination, TrainMobility, TrainEndurance, TrainPower* which are related to each other. Every exercise is linked semantically to a previous and a next exercise while every exercise has a degree of difficulty (e.g. low, medium, high, very high). Moreover, we create in the SMW for every exercise a capability, which is trained through the exercise. The exercises train the following capabilities e.g. *coordination, mobility, power, endurance*. In the next step, we generate a simulated patient with the following characteristics: *gender: male*, *age: 68*, *history of illness: arteriosclerosis*, *socialization: active, impairments: decreased coordination, endurance, power* and *mobility*, current health status is at the beginning of the evaluation set to *low*. Moreover, we create virtual sensors, which provide the information about the current user activity. Every activity is related to and defined by certain sensor state changes. The virtual sensors are connected via the WoT server. The virtual coach subscribes for these sensors in order to receive state changes. In the first step, the virtual coach computes for every exercise the optimal action path for accomplishing an exercise. The computation and creation of exercise policies lasts in average for every exercise 6 seconds. Hence, for the four exercises the virtual coach requires 24 seconds. In order to compute the average duration, we executed ten policy computation runs. In the next step, the virtual patient starts with the coordination exercise. *ActivitySensor1* reports the current state in the exercise representation. According to the state, the agent retrieves the proposed policy rule. Regarding the sent exercise states, the agent deduces correctly the appropriate actions for performing the exercise. After the agent deduces the goal state, it requests correctly the next related exercise, which is in this case the endurance exercise. This process steps are continuing until every exercise is successfully accomplished. Additionally, the virtual coach deduces and reports after every accomplished exercise, the related and trained capability. An implemented coaching avatar leads the patient through every exercise step, depending on the patient's accomplishments. The PoC shows that the virtual coaching system works as expected by us. However, we also plan to evaluate the virtual coach with the target users (e.g. physicians, physiotherapists, cardiological and neurological patients). Maybe some improvements may be required in order to simplify the usability of the virtual coaching system. We also require some strategies regarding unexpected situations. For instance, it has to be defined a strategy for situations where the patient interrupts the exercise without finishing it. Moreover, motivating strategies have to be considered in order to assure the success of the patient's rehabilitation.

6 Conclusion

The presented work has illustrated by means of an example use case the requirements of a virtual coaching system for a continuing rehabilitation at home. It has been discussed, that an IoT infrastructure is necessary to monitor the patient activities in order to improve the recovery progress of patients after a hospital stay. Therefore, the presented work provides a semantic representation framework for virtual medical exercises as well as for IoT devices. Based on this semantic representation, RL is applied in order to compute optimal exercise action paths. This is required for enabling the virtual coach to control the conduction of exercise steps by the patient. The provided *SemanticStateChart* extension supports the physicians in the creation of exercises upon which the virtual coach is enabled to provide optimal exercise policies in order to coach the patient adequate to his/her achieved capabilities and exercise level progress. However, clinical pathways are usually more complex and require more representational knowledge than the presented one. For instance, the virtual coach requires to consider also patient profiles (e.g. medical records, rehabilitation history), as well as the close interaction with the responsible physicians in order to allow the physicians to adjust the treatment of the patient. Moreover, the introduced avatar, requires to be personalized in order to provide an individual coaching experience. It is planned to transfer the presented approach also to other tasks and domains in order to show the generalizability of the presented approach.

Acknowledgement

This work is supported by the by the German Federal Ministry of Education and Research (BMBF) under the AICASys project.

References

1. Alexandrou, D., Xenikoudakis, F., Mentzas, G.: Sempath: Semantic adaptive and personalized clinical pathways

2. Bailoni, T., Dragoni, M., Eccher, C., Guerini, M., Maimone, R.: Perkapp: A context aware motivational system for healthier lifestyles. In: 2016 IEEE International Smart Cities Conference (ISC2). pp. 1–4 (Sept 2016)
3. Dragoni, M., Bailoni, T., Eccher, C., Guerini, M., Maimone, R.: A semantic-enabled platform for supporting healthy lifestyles. In: Proceedings of the Symposium on Applied Computing. pp. 315–322. SAC '17, ACM, New York, NY, USA (2017), http://doi.acm.org/10.1145/3019612.3019835
4. Laleci, G.B., Dogac, A., Olduz, M., Tasyurt, I., Yuksel, M., Okcan, A.: SAPHIRE: A Multi-Agent System for Remote Healthcare Monitoring through Computerized Clinical Guidelines, pp. 25–44. Birkhäuser Basel, Basel (2008), http://dx.doi.org/10.1007/978-3-7643-8547-7_3
5. Li, W., Liu, K., Li, S., Yang, H.: A semiotic multi-agent modeling approach for clinical pathway management. JCP 5(2), 266–273 (2010), http://dblp.uni-trier.de/db/journals/jcp/jcp5.html#LiLLY10
6. Li, W., Liu, K., Yang, H., Yu, C.: Integrated clinical pathway management for medical quality improvement – based on a semiotically inspired systems architecture. European Journal of Information Systems 23(4), 400–417 (2014), http://dx.doi.org/10.1057/ejis.2013.9
7. Shen, Y., Colloc, J., Jacquet-Andrieu, A., Lei, K.: Emerging medical informatics with case-based reasoning for aiding clinical decision in multi-agent system. Journal of Biomedical Informatics 56, 307 – 317 (2015), http://www.sciencedirect.com/science/article/pii/S1532046415001227
8. Sutton, R.S., Barto, A.G.: Introduction to Reinforcement Learning. MIT Press, Cambridge, MA, USA, 1st edn. (1998)
9. Tehrani, J., Liu, K., Michell, V.: Ontology modeling for generation of clinical pathways. Journal of Industrial Engineering and Management 5(2), 442–456 (2012)
10. qiong Wang, H., shu Zhou, T., li Tian, L., ming Qian, Y., song Li, J.: Creating hospital-specific customized clinical pathways by applying semantic reasoning to clinical data. Journal of Biomedical Informatics 52, 354 – 363 (2014), http://www.sciencedirect.com/science/article/pii/S1532046414001804, special Section: Methods in Clinical Research Informatics